

REPORT DOCUMENTATION PAGE					<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>						
1. REPORT DATE (DD-MM-YYYY) 1/1/0001		2. REPORT TYPE Journal Article - Journal Article			3. DATES COVERED (From - To) 1/1/0001 - 1/1/0001	
4. TITLE AND SUBTITLE Two Invariants of Human-Swarm Interaction				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AFRL Information Directorate Rome NY 242					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL Information Directorate Rome 32 242					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT 1 1/1/0001 12:00:00 AM						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT U	b. ABSTRACT U	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)	



Two Invariants of Human-Swarm Interaction

Daniel S. Brown

AFRL Information Directorate

and

Michael A. Goodrich, Shin-Young Jung, and Sean Kerman

Brigham Young University

The search for invariants is one of the fundamental aims of scientific endeavors. These invariants, such as Newton's laws of motion, allow us to model and predict the behavior of systems across many different problems. In the nascent field of Human-Swarm Interaction (HSI), a systematic identification of fundamental invariants is still lacking. Discovering and formalizing these invariants will provide a foundation for developing, and better understanding, effective methods for HSI. We propose two invariants underlying HSI for geometric-based swarms: (1) *collective state* is the fundamental percept associated with a bio-inspired swarm and (2) a human's ability to influence and understand the collective state of a swarm is determined by the balance between the *span* and *persistence*. We provide evidence of these invariants by synthesizing much of our previous work in the area of HSI with several new results, including a novel user study where users manage multiple swarms simultaneously. We also discuss how these invariants can be applied to enable more efficient and successful teaming between humans and bio-inspired collectives, and identify several promising directions for future research into the invariants of HSI.

Keywords: Human-robot interaction, bio-inspired swarms, recognition, human-swarm interaction, invariants, collective state, span and persistence, fan-out

1. Introduction

In the first sentence of his great paper *Invariants of Human Behavior*, Herbert Simon wrote, "The fundamental goal of science is to find invariants, such as conservation of mass and energy and the speed of light in physics. In much of science the invariants are neither as general nor as "invariant" as these classical laws" (Simon, 1990)[page 1]. In other words, science is built on invariants, concepts or properties of systems that stay the same across many different problems. As stated in Gleick's introduction to Feynman's classic lecture series (Feynman, 1964), the great insight in Feynman's lectures is that he "gathers [the] common features [of nature's laws] into one broad principle of *invariance*" [emphasis added]¹. Indeed, building a model of a phenomenon that has predictive

¹This quote is from Gleick's introduction to the 1994 Modern Library Edition of the book.

Authors retain copyright and grant the Journal of Human-Robot Interaction right of first publication with the work simultaneously licensed under a Creative Commons Attribution License that allows others to share the work with an acknowledgement of the work's authorship and initial publication in this journal.

power is precisely the process of identifying invariant properties that can be represented in such a way that they can be used in other problems.

This paper attempts to identify a set of invariants for designing systems in which a single human is managing a bio-inspired collective of robots. We generically refer to this collective as a *swarm*, but note at the onset that this term can mean both a specific form that a collective can take (e.g. a swarm of gnats (Couzin, Krause, James, Ruxton, & Franks, 2002)), as well as the label for a set of collective forms (e.g. flocks, tori, and the specific form of a swarm (Couzin et al., 2002)). We acknowledge that our list of invariants is likely to be incomplete and that future research will likely yield a list that makes some of our invariants irrelevant. Nevertheless, we argue that identifying a set of invariants in human-swarm interaction (HSI) is a useful contribution to the science of HSI. Furthermore, we seek to present sound arguments and evidence for each invariant that we propose.

This search for invariants is not new to the field of human-robot interaction. For example, the usefulness of the fan-out model (Crandall, Goodrich, Olsen Jr, & Nielsen, 2005; Olsen Jr & Wood, 2004) is that it encodes a type of invariance principle: you can't spend more time interacting with other vehicles than the neglect tolerance of any individual vehicle allows. Interestingly, neglect tolerance is a consequence of another invariance principle, namely the 2nd law of thermodynamics that says entropy increases, which implies that performance of a robot working autonomously degrades over time. Because fan-out encodes an invariant, it has specific predictive capacities, and most of the citations to fan-out identify when systems do or do not satisfy the predictions of the model.

The invariants that we propose in this paper are:

- The fundamental percept for human interaction with a swarm is the *collective state* of the swarm. Many swarms and collectives can be modeled as dynamic systems, and many of these systems exhibit attractors as defined by dynamic system theory. These attractors abstract from individual behavior to group behavior. When swarms yield such attractors, it is natural for humans to think about and manage the swarms using these attractors. In other words, the attractors become the fundamental percept of human interaction.
- There is a fundamental trade-off between *span*, the number of agents with whom the human interacts, and *persistence*, the duration of the interaction. If you can control all the agents at once you can make the group change quickly; however, if you can only control one agent in the group, then you may either have to persist a long time to effect change, or it may be impossible to effect that change at all.

This work specifically focuses on the planar geometry of swarms following simple rules of attraction, orientation, and repulsion. However, we believe that the two invariants proposed in this work apply to a larger class of swarm systems described by dynamic systems. This class includes colony behaviors that often have formal attractors such as nest selection (Nevai & Passino, 2010) and many of the collective animal behaviors described by Sumpter (Sumpter, 2010). Many interesting applications of swarms involve search (Bashyal & Venayagamoorthy, 2008), transportation (Stilwell & Bay, 1993; Rubenstein et al., 2013), and assembly (Werfel, Petersen, & Nagpal, 2014). While some of these swarm behaviors may be difficult to describe using dynamical systems, most involve simple rules of behavior followed by interacting agents that result in global, collective behaviors. Thus, we believe that the ideas presented in this paper can still provide useful insights and analogues when determining invariants for human interaction with more advanced swarm behaviors.

Before continuing we give a brief outline of the rest of the paper. We first briefly motivate the relevance of human-swarm interaction, then discuss related work. We then proceed to discuss the first invariant: collective state. The implications of this invariant are presented and subsequent sections provide evidence for each of these implications. We then discuss invariant two: span and persistence. We note that most of the results discussed in this paper are pertinent to both invariants. Thus, while the bulk of this paper is spent discussing the first invariant, most of the results discussed




Figure 1. A torus and flock formed by our model. Agent headings are represented by straight lines.

in these early sections also apply to the second invariant, and will be referenced later where appropriate. We conclude with a discussion of how these two invariants could be applied to enable more effective HRI and identify several promising directions for future research into the invariants of HSI.

In arguing for these invariants, we synthesize work that we have previously published with new work that has not been published before. We will be clear when we are referencing or reviewing previous work.

2. Human-Swarm Interaction

One of the problems with manned-unmanned teaming is that fan-out (Crandall et al., 2005; Olsen Jr & Wood, 2004; Goodrich, 2010), the number of robots a single operator can control, is limited by the twin pressures of the “bandwidth” of human attention (Pashler, 1998) and the capacity of human working memory (Baddeley, 1986). When a problem requires a single human to control large groups of autonomous agents, often the human cannot individually manage and control each agent, but instead needs to be able to manage the collective as a group. Thus, we need to allow for some agent autonomy, but keep agents simple to reduce control workload and overhead, keep costs down, provide systems resilient to failures, and calibrate trust—all common important aspects of the broad field of human factors (Lee & See, 2003; Wickens & Hollands, 2000; Sheridan, 2002).

Biological swarms provide a way to address this one-to-many goal. Bio-inspired swarms of robots are robust, distributed collectives that achieve intelligent global behaviors from a combination of simple local behaviors. Some multirobot scenarios, such as oil spill recovery or plume tracing, are amenable to having a human interact with a large swarm of robots to solve the problem (Jung & Goodrich, 2013; Spears, Spears, Heil, Kerr, & Hettiarachchi, 2005). Among the most simple behaviors that can be formed by basic self-propelled agents are a flock and a torus (see Figure 1). These are found in nature and are also fundamental attractors of simple dynamical systems with robot-realistic dynamics.

Many bio-inspired collectives have multiple locally stable group behaviors. These behaviors are often attractors of the underlying swarm dynamics. Because these behaviors are stable, they allow a human to interact with the swarm by injecting small amounts of information to shape the collective’s behavior, without causing it to leave a local attractor. A human can alternatively perturb the system to “bump” it into a different attractor. This provides a method for dynamic control of collective behaviors that are also stable or that switch between stable points, providing a user with the ability to neglect one swarm and then switch control to another swarm. Influencing the collective pattern of the swarm using only a small number of agents has the potential to reduce the communication and/or cognitive bandwidth needed to model the state of the swarm and control its actions, thus enabling more operator-friendly control of the swarm by allowing an increase in fan-out and a decrease in

workload.

3. Related Work

Bio-inspired swarm models have been explored by researchers in a wide variety of fields including computer science, engineering, physics, and biology. These models are typically capable of either flocking (Reynolds, 1987; Vicsek, Czirók, Ben-Jacob, Cohen, & Shochet, 1995; Jadbabaie, Lin, & Morse, 2003; Olfati-Saber, 2006) or cyclic behavior (Marshall, Broucke, & Francis, 2004; Levine, Rappel, & Cohen, 2000), and in some cases can exhibit multiple group behaviors depending on the model parameters used (Couzin et al., 2002; Strömbom, 2011; Romero, Forgoston, & Schwartz, 2011).

Many researchers have also investigated influencing the behavior of swarms. Couzin et al. (Couzin, Krause, Franks, & Levin, 2005) and Conradt et al. (Conradt, Krause, Couzin, & Roper, 2009) explore leading a flock with a small number of informed agents. Couzin *et al.* show that their method of leading a flock scales well as group size increases. Additional work (Olfati-Saber, 2006) uses global information and communication with local neighbors to form a robust flock and proves that a leader agent can lead the group through global information. Su et al. (Su, Wang, & Lin, 2009) eliminate the need for global information and show that only a subset of informed individuals are needed to direct a flock. Jadbabaie (Jadbabaie et al., 2003) provides mathematical results on the convergence of the Vicsek flocking model (Vicsek et al., 1995) to a single group direction as well as convergence conditions for the group to converge in the limit to a single leader’s direction.

Previous work on human-robot interaction has contributed approaches for estimating fan-out using experiments from only a single robot (Crandall et al., 2005). Other work has estimated fan-out using large simulation studies involving several sizes of robot teams (Olsen Jr & Wood, 2004; Olsen Jr, Wood, & Turner, 2004). Recent work by Glas and Zheng has studied fan-out with real robots in social settings such as navigation and shopping (Glas, Kanda, Ishiguro, & Hagita, 2012; Zheng, Glas, Kanda, Ishiguro, & Hagita, 2013b, 2013a). Pourmehri et al. examine group-level abstracted multi-robot control to form dynamically-sized teams to increase fan-out (Pourmehri, Monajjemi, Vaughan, & Mori, 2013). We differ from this set of previous work because we are interested in managing the group as a whole rather than as individuals. This means that the group is treated as a single entity, allowing us to use fan-out as an indirect measure of workload by asking how many swarm units a single human can control.

Research in the emerging field of Human-Swarm Interactions (HSI) has often used user studies to investigate workload and performance. However, these studies have not focused on managing swarm attractors and have not tried to predict workload and performance. Recent work has compared human-swarm interactions using different leadership and control methods (Pendleton & Goodrich, 2013; Kolling, Sycara, Nunnally, & Lewis, 2013). Nunnally et al. explore bandwidth constraints on swarm-to-human communications, but assume that the human can communicate with all of the agents in the swarm (Nunnally et al., 2012). Walker et al. investigate communication latency between a human operator and a swarm and demonstrate *neglect benevolence*, the importance of allowing the emergent behaviors of a swarm to stabilize before giving new commands (Walker et al., 2012). Based on recent research regarding neglect benevolence (Nagavallik, Luo, Chakraborty, & Sycara, 2014; Nagavalli, Chien, Lewis, Chakraborty, & Katia, 2015), we believe this may be another invariant of HSI; however, we do not focus on neglect benevolence in this paper.

Work by Cummings has investigated supervisory control of swarms of UAVs; however, this work differs from our work because the systems used are not described by the dynamical models that are critical to our invariants (Cummings, 2004). Recent results have also made attempts to directly measure how a human perceives a swarm; initial results, though encouraging, still need further validation (Manning, Harriott, Hayes, Adams, & Seiffert, 2015; Harriott, Seiffert, Hayes, &

Adams, 2014; Seiffert, Hayes, Harriott, & Adams, 2015).

4. Invariant One: Collective State

In this section, we argue that the fundamental percept of HSI is the state of the swarm. This is a principle of abstraction (Goodrich & Mercer, 2012); in this case, the collective state abstracts the highly complex local interactions of individual agents. Indeed, understanding the relationships between collective state and the behavior of individual agents is a relatively recent contribution of science (Herbert-Read et al., 2011; Sumpter, 2010; Seely, 2010; Levine et al., 2000; Couzin & Krause, 2003; Ballerini et al., 2008; Nevai & Passino, 2010; Reynolds, 1987). Claiming that the collective state is the fundamental percept requires that the human is able to perceive, understand, and influence the abstracted collective state, even when it is computationally prohibitive to understand the behavior of individual agents. To understand this claim better, consider the following observations from the literature.

One reason for studying swarms in both biology and robotics is the emergence of repeatable patterns; as stated by Sumpter, “Linking different levels of organization involves the study of collective phenomena: phenomena in which repeated interactions among many individuals produce patterns on a scale larger than themselves” (Sumpter, 2010, p.1). Indeed, much of Sumpter’s book evaluates simplified mathematical models of individual interactions that produce repeatable collective patterns (Sumpter, 2010). Couzin’s seminal work in this area identifies four collective patterns: a swarm, a torus, and two forms of what we call a flock and what he calls parallel groups (Couzin et al., 2002). Furthermore, Couzin makes the observation that the collective patterns act as a form of memory, suggesting that a collective pattern is a type of state in some sort of computational function performed by the group. Similarly, Seeley’s fascinating work on nest selection in honeybees emphasizes transitions from one collective pattern to another in response to the need to “compute” or solve problems that the group faces (Seely, 2010). From a control theory perspective, work in flocking and foraging relies on the robustness of these collective patterns, and much work has been dedicated to deriving decision rules for individual agents that, when combined, produce reliable patterns (Olfati-Saber, 2006; Jadbabaie et al., 2003; Mesbahi & Egerstedt, 2010).

The presence of the repeatable collective patterns says little about how a human might influence these collective patterns. Initial work on influencing the state of the swarm focused either on broadcasting new parameters to the all agents in the collective, inducing them to change collective state (Nunnally et al., 2012; Kolling et al., 2013; Kira & Potter, 2010), ad hoc approaches for influencing subgroups (Bashyal & Benayagamoorthy, 2008), or on designating a specific leader agent (which could be a collocated human (Alboul, Saez-Pons, & Penders, 2008)) that could operate under human influence while other agents maintained the appropriate pattern given changes in leader behavior (Jadbabaie et al., 2003; Vaughn, Sumpter, Henderson, Frost, & Cameron, 1998; Lien, Bayazit, Sowell, Rodriguez, & Amato, 2004). The limitation to these approaches, of course, is that it puts the human into the role of a centralized controller. This may be desirable in some problems, but seems to violate the benefit from having a robust swarm, where each agent is contributing to the collective behavior.

By claiming that an attractor of a dynamic system forms a collective state that is the fundamental percept of human-swarm interaction, we are essentially claiming (A) that a human should be able to manage swarms by managing attractors. Such a claim implies that there exist useful robot swarm models that have (B) collective states that are stable, recognizable attractors. It also implies that a human should be able to (C) switch between these collective states, (D) identify a collective state in the absence of complete information from all the agents, and (E) modify the behavior of a particular collective state. We provide evidence for the claim by providing evidence for each of these

implications, recognizing the logical limitation of trying to prove A using the observation that

$$A \Rightarrow B \wedge C \wedge D \wedge E \quad (1)$$

and generating evidence for B , C , D , and E . Within the constraints of this limitation, we begin by providing evidence for each. We begin with implication B .

4.1 Implication B : Collective States of a Robotic Swarm

A solution to the challenge of allowing a human to influence the collective state without becoming a central point of control is to treat the collective state of a set of distributed agents as a stable attractor of the individual agent dynamics. Such stable attractors have been identified in models of honeybee behavior (Nevai & Passino, 2010) and artificial agent societies (Strömbom, 2011; Romero et al., 2011); the stable attractor abstraction is the foundation for the remainder of this section. We have made extensive use of this in our prior work in HSI and will review some of this work in the next subsection while presenting a new result that provides evidence that the flock behavior is a stable attractor of the swarm model considered in this section.

The main idea behind managing a swarm by managing its attractors is to find a set of parameters that allow multiple collective states to emerge. In control theoretic terms, this means that we are looking for parameters that yield locally stable attractors (Nevai & Passino, 2010). The local stability allows either environmental input obtained by the agents or influence from a human to guide the collective from one basin of attraction to another, inducing a change in state.

4.1.1 A Useful Swarm Model The experiments and simulations in this paper build upon on a previously published model of swarming (Kerman, Brown, & Goodrich, 2012). We have chosen to focus on this model because it has two collective behaviors (a flock and a torus), is similar to many biological models of swarming behavior, and has dynamics similar to those of actual robots. This model consists of a set of N agents, modeled as planar unicycles, with dynamics for agent i given by

$$\dot{x}_i = s \cdot \cos(\theta_i), \quad \dot{y}_i = s \cdot \sin(\theta_i), \quad \dot{\theta}_i = \omega_i \quad (2)$$

where $\theta_i \in [-\pi, \pi]$ is the agent's angular heading, s is the constant agent speed, and ω_i is the agent's angular velocity. For simplicity we use the notation $c_i = [x_i, y_i]^T \in \mathbb{R}^2$ and $v_i(t) = [\cos(\theta_i(t)), \sin(\theta_i(t))]^T$ to represent agent i 's position and velocity, respectively. Similar to many bio-inspired models, we assume that agents can sense neighboring agents' positions and headings, but cannot explicitly communicate or coordinate. We use a random neighbor model where agent j is visible to agent i at time t according to a Bernoulli random variable with parameter $p_{ij}(t) = \min(1, 1/d_{ij}(t))$, where $d_{ij}(t)$ is the Euclidean distance between agents i and j at time t . We chose this random neighbor model because it is similar to the random neighbor model used in (Bode, Franks, & Wood, 2010) which replicated field observations of starlings (Ballerini et al., 2008), and is also relevant for noisy, distance-limited robot sensors.

Similar to the Couzin model of biological swarms (Couzin et al., 2002) and the Reynold model of synthetic agents (Reynolds, 1987), agents in our model react to neighbors within three different zones: repulsion, orientation, and attraction. The neighbors in these zones are determined by

$$n_i^r = \{j : \|c_i - c_j\| \leq R_r, a_{ij} = 1\} \quad (3)$$

$$n_i^o = \{j : \|c_i - c_j\| \leq R_o, a_{ij} = 1\} \quad (4)$$

$$n_i^a = \{j : a_{ij} = 1\} \quad (5)$$

where n_i^r , n_i^o , and n_i^a are the sets of agent i 's neighbors in the regions of repulsion, orientation, and attraction, respectively. Throughout this paper, $\|\cdot\|$ is used to denote the standard Euclidean

norm. The parameters R_r and R_o are the associated radii of repulsion and orientation. The angular velocity ω_i is determined by first computing the repulsion, orientation, and attraction vectors

$$u_i^r = -\sum_{n_i^r} \frac{c_j - c_i}{\|c_j - c_i\|^2}, \quad u_i^o = \frac{v_i + \sum_{n_i^o} v_j}{\|v_i + \sum_{n_i^o} v_j\|}, \quad u_i^a = \frac{\sum_{n_i^a} (c_j - c_i)}{\|\sum_{n_i^a} (c_j - c_i)\|}. \quad (6)$$

Next, the desired heading vector u_i is computed as

$$u_i = u_i^r + u_i^o + u_i^a. \quad (7)$$

Finally, angular velocity, ω_i , is computed as $\omega_i = k(\text{atan2}(u_i^y, u_i^x) - \theta_i)$ where k is a positive gain and $\text{atan2}(u_i^y, u_i^x)$ is the two argument variation of the arc-tangent that places the angle in the correct quadrant by considering the signs of the y and x components of u_i . Because we limit $(\text{atan2}(u_i^y, u_i^x) - \theta_i)$ to the interval $[-\pi, \pi]$, the magnitude of ω_i is bounded by $k\pi$, matching the Dubins curve model (Dubins, 1957) often used for single and multiple UAV path planning (Chitsaz & LaValle, 2007; Ding, Rahmani, & Egerstedt, 2010).

4.1.2 Flock and Torus Attractors The model produces two emergent behaviors: a torus and a flock. Snap shots of these behaviors are shown in Figure 1. A torus is characterized by a relatively stationary group centroid, and either a clockwise or counterclockwise rotation. A swarm in the torus formation could be used for perimeter monitoring, or as a fixed-wing UAV loitering command. A flock is characterized by a moving centroid with all of the agents moving in the same general direction. A swarm in the flock formation is ideal for transporting the swarm quickly from one location to another and could be used for search or tracking.

Kerman et al. analyzed the torus attractor and showed that it is a fundamental attractor of the swarm model presented in (Kerman et al., 2012). This paper contributes the following argument that the flock behavior is fundamentally caused by the orientation dynamics in our model.

Consider a system with only orientation dynamics ($u_i = u_i^o$ in Equation (7)). The angular velocity $\dot{\theta}_i$ of agent i becomes

$$\dot{\theta}_i = k \left(\text{atan2} \left(v_i^y + \sum_j v_j^y, v_i^x + \sum_j v_j^x \right) - \theta_i \right). \quad (8)$$

We can approximate the right-hand side of (8) using the following form of angular averaging

$$\begin{aligned} \dot{\theta}_i &= k \left(\frac{\theta_i + \sum_j \theta_j}{n_i + 1} - \theta_i \right) \\ &= k \left(\frac{(\theta_i + \sum_j \theta_j) - (n_i + 1)\theta_i}{n_i + 1} \right) \\ &= \frac{k}{n_i + 1} \sum_j (\theta_j - \theta_i) \end{aligned} \quad (9)$$

where n_i is the number of neighbors of agent i . We can define the *orientation graph* where each node is a member of the swarm and where an edges exist between members of the swarm that are neighbors, i.e., those members of the swarm that interact. When the underlying orientation graph is connected, (9) is known to cause all agents to converge to a common heading (Ren & Beard, 2007; Mesbahi & Egerstedt, 2010). Simulation results show that even when the orientation graph dynamically switches, agents converge to a common heading. Thus, we argue that the flock is




Figure 2. Abstracting individual agent behaviors and focusing on managing collective attractors allows us to view the swarm as a finite state machine where limited perturbations can cause state transitions.

fundamentally caused by the orientation dynamics and is an attractor of the swarm model. The attraction and repulsion dynamics serve to keep the flock cohesive and prevent collisions. This analysis is essential because it means that we can treat a flock as a cohesive unit, allowing a human to influence collective flock behaviors without needing to invoke centralized control over each agent. Additionally, because the flock behavior is a stable attractor, a human operator can temporarily neglect the flock to accomplish other tasks without the flock decohering or drastically changing its behavior.

4.2 Implication *C*: Switching Between Attractors

Recall that implication *C* from Equation (1) says that one form of human influence is to induce the collective to change states, something that is possible if the human perceives and interacts with the swarm as a collective unit. Having established that there exists a useful model of robot dynamics and inter-robot interactions that yields multiple stable attractors—implication *B* from Equation (1)—we turn to implication *C*. An important constraint is that we do not want the human to simply take control of all agents: this makes the human a single point of failure and fails to scale when the number of agents grows and bandwidth is limited. Instead we focus on the case where the human can only directly influence a subset of the swarm.

Abstracting to the attractors from the previous section creates a finite state machine with two states: a flock and a torus (see Figure 2). A two-state state machine is certainly within human capacity, and is much more compatible with human limitations than managing all the possible configurations of a multi-agent complex dynamic system. The goal for the remainder of this section is to construct feasible mechanisms that a human could use to manage switches between these collective attractor-states. Feasibility is constructed through citations to our previous work with human participants managing swarms, and to a simple Oz of Wizard study where the Wizard is a simulated human. We acknowledge that simulated robots with simulated humans is a relatively weak form of evidence, but taken with the other forms of evidence in this paper we argue that this “simulated on simulated” approach yields useful information.

4.2.1 Switching between attractors using stakeholders Throughout this paper we assume that the human operator has limited bandwidth and therefore must influence the entire swarm by sending a command to only a subset of agents in the swarm. In Section 6, we will explicitly discuss this assumption and distill this into the second invariant. We refer to the subset assumption as the *span property*, and note that the span property applies to situations where only a subset of the agents are capable of communicating to a human due to power or equipment constraints, as well as sit-

uations involving contested or noisy communication environments. In this section, we adopt the terminology in (Kerman et al., 2012) and refer to the subset of agents that the human can influence as *stakeholders*. Stakeholders are influenced by both the human and by other agents.

In this subsection, we review part of the analysis and modeling performed in (Brown, Kerman, & Goodrich, 2014) and show that both flocks and toruses can be led using two different methods: attraction and orientation. Stakeholders that are led by attraction have the desired direction $u_i = u_i^{sa} + u_i^o + u_i^r$ where

$$u_i^{sa} = \frac{\rho \hat{q}_i + (1 - \rho)u_i^a}{\|\rho \hat{q}_i + (1 - \rho)u_i^a\|_2} \quad \text{and} \quad \hat{q}_i = \frac{q - c_i}{\|q - c_i\|_2}. \quad (10)$$

The vector $q \in \mathbb{R}^2$ is the human-generated waypoint, $\rho \in [0, 1]$ is priority parameter, and u_i^a, u_i^o, u_i^r are the usual attraction, orientation, and repulsion influences described previously in Equation (6). This causes stakeholders to continue to orient with their neighbors while adjusting their positions to be closer to the human-specified waypoint.

Stakeholders that are led by orientation have the desired direction $u_i = u_i^{so} + u_i^a + u_i^r$ where

$$u_i^{so} = \frac{\rho \hat{q}_i + (1 - \rho)u_i^o}{\|\rho \hat{q}_i + (1 - \rho)u_i^o\|_2}. \quad (11)$$

This causes stakeholders to continue to be attracted to their neighbors' positions while adjusting their headings towards the human-specified waypoint. We use the notation N_{stk} to denote the number of stakeholders.

We assume for the sake of simplicity that an operator may only command absolute positions, and commands have no persistence. While this may not be true for many swarm systems, it allows us to study human interaction with swarms made up of extremely simple agents following classic "boids"-like rules (Reynolds, 1987), where each agent changes its behavior at each time step by reacting only to current percepts from its neighbors and, in our case, from the current operator input. Future work should address this assumption and investigate what happens when agents have memory or when other operator inputs are available.

4.2.2 Oz of Wizard Study The results in this subsection were first presented in (Brown, Kerman, & Goodrich, 2014). The goal of this study was to investigate the stability of switching from flock-to-torus and torus-to-flock, under a range of parameter values, using the two different leadership strategies: leadership by attraction and leadership by orientation. To determine how best to switch the swarm from one group type to the other, we ran simulations with values of N_{stk} (the number of stakeholders) ranging from 10 to 100 in 10 step increments and values of ρ (the relative priority of the human input) ranging from 0.1 to 1 in 0.1 step increments. Both orientation and attraction leadership strategies were tested using a constant waypoint command. These methods systematically explore the bounds of what an idealized human, which we denote as an OoWiz after the Oz of Wizard approach (Steinfeld, Jenkins, & Scassellati, 2009), can do to switch between attractors.

Ten trials were performed for each N_{stk} and ρ pair, and for each leadership strategy. To initialize the simulations, agent positions were distributed randomly in a 20 unit \times 20 unit square centered at the origin. To ensure the swarm started as a flock, initial orientations were all set to 0. To ensure the swarm started as a torus, initial orientations were set to $\theta_i = \text{atan2}(c_i^x, c_i^y) + \pi/2$. After letting the group stabilize for 25 seconds, an arbitrarily large constant control input, $\hat{q} = c_g(25) + [10000, 0]^T$, was applied to the stakeholders to influence them to form a flock. We gave the swarm 200 seconds to switch group types, removed the OoWiz human influence, and gave the swarm 50 seconds to stabilize to evaluate whether the swarm would remain in the desired group type.




Figure 3. Switching from flock to torus.

4.2.3 Switching from Flock to Torus Leadership by attraction was effective for switching from a flock to a torus. Figure 3(a) shows the percentage that switched when under the influence of the OoWiz human. Figure 3(b) shows the percentage that remained in the torus formation after the OoWiz human input was removed. As can be seen, for sufficiently high N_{stk} and ρ , the group successfully switched. However, there is a noticeable drop in the number of simulations that switched and remained for $\rho = 1$ and high values of N_{stk} . We investigated this and found that because the OoWiz human was explicitly controlling the attraction dynamics, the agents formed a flock-like structure that circled around the reference input. Because the agents never spread out into a full torus, when the control input was removed the group returned to a flock formation.

When leading stakeholders by orientation, we found a large discrepancy between the number of simulations that switched to a torus and the number of simulations that remained as a torus (see Figures 3(c) and 3(d)). We investigated this phenomenon and found that the human influence on the orientation dynamics often caused stakeholders to rotate in different directions around the desired torus centroid. This behavior prevents the torus from fully forming and is undesirable for actual robots because of the risk of head-on collisions. After the human influence was removed, the stakeholders were able to orient properly, causing a torus to form. This is an example of what Walker et al. refer to as *neglect benevolence* (Walker et al., 2012). Neglect benevolence says that, in some cases, the swarm must be allowed to self stabilize before receiving a new command from the human.

4.2.4 Switching from Torus to Flock When switching from a torus to a flock (see Figure 4), leading stakeholders by orientation worked much better than leading stakeholders by attraction. We examined these results and found that leading stakeholders by attraction was successful in causing the agents to switch from a torus to a flock. However, the attraction input \hat{q} caused the stakeholders to slowly pull away from the rest of the group causing the flock to elongate. Thus, when the human influence was removed, the flock was unstable and usually reformed a torus. Figures 4(a) and 4(b) show that except for limited areas of the parameter space, simulations that switched to a flock usually switched back to a torus. Using leadership by orientation eliminated this phenomenon and caused the agents to form a less elongated flock that remained stable after human influence was removed (see figures 4(c) and 4(d)).

4.2.5 Discussion We have shown that we can switch between attractors.² Switching between collective states or attractors is useful when we want to have the swarm change its collective behavior. Having stable attractor-states is critical if the swarm is to operate in environments where agent failure or malfunction may occur, as it provides robustness to small errors and perturbations. Having

²Further improvements have been shown for switching collective states using the principle of quorum sensing (Brown, Kerman, & Goodrich, 2014; Sumpter, 2010).




Figure 4. Switching from torus to flock.

multiple stable collective states allows the human operator to manage a swarm at the attractor level. This reduces the workload required to change a swarm’s behavior and also reduces operator workload by providing the opportunity to occasionally neglect the swarm while having some guarantee of the collective behavior remaining stable. We also note that because these results involve fundamental system dynamics, they are not only useful for human control, but also provide useful information for the design of an algorithmic controller.

4.3 Implication *D*: Identifying Collective State

Recall that implication *D* from Equation (1) says that if collective state is a fundamental percept, then humans should be able to identify which state the system is in by observing a subset of agents (the span property) over a period of time. We refer to observing or influencing agents over a period of time as the *persistence property* and note that the results from the previous section implicitly rely on the persistence property; collective state was measured after a fixed time interval of human influence and then human neglect.

Returning to implication *D*, subjective evaluation indicates that it is easy for a human to observe a subset of agents for a relatively small period of time and identify the collective state, flock or torus. As a simple example consider the trajectories shown in figures 5 and 6. They show a subsample of agents from a flock and a torus where we have captured the trajectories of the agents for a certain amount of time steps. Note that when observing 20% of the swarm, we can observe for a small period of time (2 time steps) and obtain an accurate estimate of the global behavior. Alternatively, we can watch only 2% of the swarm for a longer period of time (20 time steps) and also obtain an accurate estimate. Thus, given the right span (number of observed agents) and persistence (length of observation), with high probability, you can tell whether the collective state is flock or torus.

Adding to this subjective evidence, we constructed an extremely simple classifier, a naive Bayes classifier to be precise, and varied how many agents it could observe and for how long (Brown & Goodrich, 2014). Even in the presence of noise, the classifier works extremely well over a wide range of span and persistence (see Section 6 for a discussion of those terms). Shown in Figure 7 are the results where we sampled agents’ angular velocity and numbers of neighbors and used these features to classify whether the swarm was a clockwise torus, a counterclockwise torus, or a flock. With high accuracy, we can predict the global state of the swarm from this limited information. We will return to this figure in Section 6 when we discuss span and persistence in more detail.

4.4 Implication *E*: Managing a Torus

Recall that implication *E* from Equation (1) says that if the collective state is a fundamental percept, then people should be able to manage the behavior of that collective state. This is easily understood by considering a flock. The flock state is characterized by agents aligning with each other and moving in a coherent direction, but the direction of the flock should be responsive to the needs of the



Figure 5. Trajectories of 2% of a torus and a flock observed for 20 time steps.




Figure 6. Trajectories of 20% of a torus and a flock observed for 2 time steps.

collective. Thus, *another form of human influence is to “nudge” the pattern of collective behavior to perform some type of useful function.* In this section we present the details of modifying the collective state of a torus. In the next section we focus on managing a flock.

The first piece of evidence in support of Implication *E* we presented in (Goodrich, Pendleton, Kerman, & Sujit, 2012). In the simulation study in that paper, we explored how a set of stakeholders could be used to move a torus. Evidence from this paper shows that selecting a few stakeholders and slowly moving them, while allowing them to stay in the torus shape, causes the torus to slowly shift its center from one location to another. In this case, the modifier of the torus is the (x, y) -location of the torus center, which can be modified using the now-familiar stakeholder form of influence.

The next piece of evidence for Implication *E* for a torus we presented in (Kerman et al., 2012). In this paper, we designed a very simple controller for a set of stakeholders in a torus. The simple controller simply caused the stakeholders to be attracted to a slow time-varying reference input. Since the maximum velocity of the center of the torus has a tight upper bound equal to $\frac{\pi}{4}$ times the speed of the agents (Jung, Brown, & Goodrich, 2013), if the reference input moves too quickly




Figure 7. Collective behavior classification accuracy as the number of agents sampled increases for different numbers of successive samples per agent. Results are for a swarm with 100 agents. We used a simple naive Bayes classifier where the features were each agent’s angular velocity and number of neighbors.

it is not surprising that the center of the torus cannot track it. However, as long as the path does not change too quickly, the attraction of stakeholders to the path causes the center of the torus to track the path. This implies that if a human slowly influences a torus by “tugging” it along via stakeholders, the human can control how the center of the torus moves in the world.

A more interesting study, and the third piece of evidence, is one where we explored how humans could alter the shape of the torus boundary, changing it from a circle to a rounded square or other rounded polygonal shapes (Jung et al., 2013). To accomplish this, we had to modify the way that the human influenced the agents in the collective, and we had to make small modifications of the agents themselves. Rather than influencing the collective with stakeholders, the humans influenced used special agents called *mediators*. A mediator is a human-controlled agent that has its own regions of attraction and repulsion; agents within the region of repulsion are repelled by the mediator and agents outside this region but within the region of attraction are attracted to the mediator. This allows mediators to draw agents near to them but not too near, and, because of the circular nature of the regions, allows the mediators to bend the paths of the agents. Very complex polygons can be approximated by placing some mediators on the inside of the torus, causing the torus to bend outward, and other mediators on the outside, causing the torus to bend inward. The problem with this approach, however, is that it requires too many mediators, and this is undesirable because controlling a lot of mediators is just as difficult as controlling a lot of agents in the swarm.

In our work in (Jung et al., 2013), we introduced the mediator idea to a torus constructed using precisely the dynamics from Couzin’s model (Couzin et al., 2002). As can be seen in Figure 8, the number of mediators required to produce even simple polygons was surprisingly high. The reason that so many mediators are required is that the agents in the torus have too much influence over each other (Brown, Jung, & Goodrich, 2014); they “see” each other across the mediators and the resulting inter-agent influence tends to overpower the influence of the mediators (Figure 9(a) left). By reducing inter-agent influence to the nearest neighbor in front of an agent (Figure 9(a) right), the mediator is better able to shape the torus (Figure 9(b)). The cost is that a small amount of memory must be added to each agent allowing it to remember the last position of its nearest neighbor, allowing the members of the swarm to follow their nearest neighbor despite occasional occlusions; without this, the torus is unstable (Jung et al., 2013).




Figure 8. High inter-agent influence (Figure 9(a) left) results in too many mediators required to form simple polygons; adapted from (Jung et al., 2013).

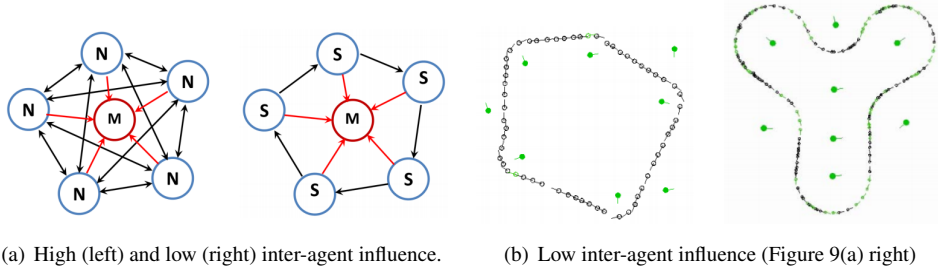


Figure 9. Decreasing inter-agent influence allows more mediator influence; adapted from (Jung & Goodrich, 2013).

The user study presented in (Jung & Goodrich, 2013) also provides evidence for implication *E* for the torus attractor. In this user study we used mediators to allow a human to control the shape of a torus formed by the modified agents mentioned in the previous paragraph. The user study required the human to surround an expanding oil spill using simulated agents. Different shapes of oil spills were used, each roughly corresponding to a different polygonal shape, and measures of efficiency and effort indicate that the human was able to shape the torus to match the different desired shapes.

4.5 Implication *E*: Managing a Flock

Recall that implication *E* from Equation (1) says that if the collective state is a fundamental percept then people should be able to manage the behavior of that collective state. In the previous section, we presented evidence for this implication with respect to the torus. In this section and the sequel, we present evidence for this implication with respect to the flock.

As in the previous section, the first piece of evidence in support of this we presented in (Goodrich et al., 2012). In the simulation study in that paper, we also explored how a set of stakeholders could be used to move a flock. Evidence from this paper shows that selecting a few stakeholders and slowly changing their headings causes the entire flock to change its direction of travel to the desired heading. In a follow-on paper, we conducted a user study in which we allowed humans to control flocks of Couzin-like agents in a foraging problem (Pendleton & Goodrich, 2013). This user study is the second piece of evidence to support Implication *E* for flocks. In the user study, the human was allowed to manage a small set of stakeholders. The human was also allowed to influence the flock by controlling either a leader “super-agent” or a predator “super-agent” that caused the

nominal agents to be attracted to or repelled by the super-agent. Leaders and stakeholders produced equivalent performance, both higher than predators, and both had lower workload than predators. This indicates that stakeholders are as easy for a human to use as a leader super-agent, but unlike the leader, which causes a torus to turn into a flock, stakeholders can be used to control a torus too.

Again, as in the previous section, the next piece of evidence for Implication *E* for flocks we presented in (Kerman et al., 2012). In this paper, stakeholders in a flock were attracted to a time-varying point. For a sufficient number of stakeholders and for a point that moves slowly enough, the stakeholders can cause the flock to track the point. This should not be surprising since the user study in (Pendleton & Goodrich, 2013) indicated that humans can cause the flock to fly to desired locations by influencing stakeholders, but it is nevertheless another piece of evidence for Implication *E*.

As the final piece of evidence that the collective state is a fundamental percept that allows a human to manage the behavior of that collective state, we present results from a new user study. Because the user study has many elements, it is more convenient to discuss it in a separate section.

5. Managing Multiple Flocks: A New User Study

The purpose of the study presented in this section is to stretch the concept of treating a swarm as a collective unit. The study does this by having the human manage multiple flocks. There are two key elements of this: if a flock is a stable attractor, a user should be able to neglect it for a period of time without it decohering; and if the flock is a collective unit, a user should be able to recognize and modify the behavior of the flock when only a portion of the flock can be seen.

The user study uses the notion of fan-out (Olsen Jr & Wood, 2004; Crandall et al., 2005), but rather than measuring fan-out as the number of homogeneous, independent robots that a human can manage, we instead measure fan-out as the number of homogeneous, independent flocks that a human can manage. We use three metrics of human workload and use these metrics to formulate a model for the fan-out of an operator managing multiple swarms in the flock formation. We show that the model provides adequate but coarse-grained approximations of workload. A consequence of this is that the model predictions can be used to design swarm systems with desired fan-outs and workloads in mind. The key reason this is possible is that we are managing attractors of a dynamical system that are naturally stable and therefore are amenable to periods of interaction and neglect.

5.1 Metrics for Human Swarm Interaction

We are interested in measuring how well a human can interact with a collective in the flock state. In section 4.1.2 we provided evidence that the flock is an attractor of the dynamic system. This is key because it allows us to abstract human influence from agent-level interactions to attractor-level interactions. To show the strength of having an attractor based model we show that the commonly used human-robot metric of fan-out can be applied to robot swarms. In the case of bio-inspired swarms, fan-out (FO) encodes the number of flocking groups that can be managed by a human. If each group is composed of N agents, then the human can manage $N \cdot FO$ agents.

Fan-out is defined as a modified version of the equation given in (Crandall et al., 2005):

$$FO = \frac{NT}{IT + OT} + 1. \quad (12)$$

This is a function of three system characteristics: interaction time (IT), neglect time (NT), and observation time (OT). *Interaction time* is the time it takes a human to reduce the swarm’s performance error below some desired threshold level. IT is an explicit encoding of persistence (duration of control input) given a specified level of span (number of agents being controlled). *Neglect time* is the time it takes for the swarm’s performance error to rise above some error threshold when the robot is left unattended by the human. NT relies on an understanding of the invariant state, noting that




Figure 10. Scenario for influencing a flock

the pattern is not likely to change given random perturbations because the system has been designed with large enough basins of attraction for the two states, but that the behavior within the state may drift given random perturbations. *Observation time* is the time required to estimate the true collective state of a swarm within a given error threshold. OT is an explicit encoding of persistence given a specified level of span, relating to the amount of time required to understand the behavior of the agents within a state.

We note that whereas IT and NT relate to the desired state of the swarm, OT relates to the actual state of the swarm. Thus after neglecting a swarm the human must observe the swarm to estimate its actual state and then interact with the swarm to get the actual state to match the desired state. We also acknowledge that in many cases observations and interactions may be interleaved and simultaneous. Thus, it may not be possible to separate IT and OT. However, in this study we choose not to explicitly model complex interactions between IT and OT. Instead, we simplify our analysis by estimating each individual quantity and then test whether the sum $IT + OT$ provides a close enough approximation such that fan-out can be predicted.

5.2 Experiment Design

We designed a user study to determine how well we can predict values for fan-out—given estimates of interaction time, neglect time, and observation time—that match actual human performance and perceived workload. Controlling the heading of a flock by influencing a subset of the swarm is a popular research problem; however, previous work has only studied this problem for a single flock (Couzin et al., 2005; Olfati-Saber, 2006; Genter, Agmon, & Stone, 2013). Our user study is the first research that involves a tracking problem where users must control the headings of four swarms simultaneously.

Our user study consisted of four scenarios where in each scenario the task given to the users was to control four different swarms to simultaneously track four different target headings. The human was only able to interact with the swarm by influencing the headings of a subset of the swarm through a movable waypoint (see Section 4.2 for details on the lead-by-attraction model). In each scenario the flock started off moving in the incorrect direction so the user had to influence the flock to get it to turn and go in the right direction. Figure 10 shows an example of this scenario. The flock starts heading in the direction θ_0 and the goal of the user is to get the flock to go in the desired direction θ^d . The target heading randomly changes at each time step according to a Wiener process that will be described below.

The user received a performance score at the end of each of four different scenarios based on how well the user was able to keep each swarm’s actual heading within a fixed neglect tolerance of

the desired heading. This score was only displayed at the end of each scenario to avoid giving the user additional information about the headings of each swarm, but still gave each user an incentive to perform the task as well as possible. Each user completed the same four scenarios, but in a random order to counterbalance any learning that might have taken place. Each of the four scenarios was designed to have a specific fan-out based on conditions involving interaction time, observation time, and neglect time. These specific fan-out values were calculated using estimates of interaction time, neglect time, and observation time. The following sections describe each of these estimates.

5.2.1 Interaction Time To obtain a model of interaction time we ran a simple case study of 8 individuals (7 male and 1 female) to collect data on human control of a flock to get it to turn to a desired heading. We generated random desired headings between 10 and 80 degrees out of phase with the actual heading of the flock and told the users to turn the flock to the desired heading as quickly as possible while keeping the swarm cohesive. Each user was given 20 scenarios and each flock comprised 100 agents. For half of the scenarios the user controlled the swarm using 40 stakeholders (corresponding to a difficult interaction) and for half of the scenarios the user controlled 70 stakeholders (corresponding to an easy interaction). We found that the average turning speed was 2.96 degrees/second for 40 stakeholders and 4.94 degrees/second for 70 stakeholders. Using the data collected from the case study we fit linear equations to the data using the MATLAB `robustfit` and obtained the following equations

$$IT_{40}(e_\theta) = 0.1723 \cdot e_\theta + 10.468 \quad (13)$$

$$IT_{70}(e_\theta) = 0.1347 \cdot e_\theta + 4.5922 \quad (14)$$

where e_θ is the error between the current and desired headings of the swarm and $IT_{40}(\cdot)$ and $IT_{70}(\cdot)$ represent the interaction time when using 40 or 70 stakeholders, respectively. We use these equations later in this paper to estimate the interaction time for a human to change a flock's heading.

We used these values to split our interaction time condition into an easy and hard parameter set. Based on user feedback, we chose to use a neglect tolerance of $\pm 5^\circ$. This made the task difficult, but not impossible. We used equations (13) and (14) to solve for the interaction time necessary to readjust the heading of swarm that is about to be off course, resulting in $IT_{40}(5^\circ) \approx 11.33$ and $IT_{70}(5^\circ) \approx 5.26$. We combined these estimates with the following estimates of observation time and neglect time to estimate fan-out.

5.2.2 Observation Time Observation time is often neglected in studies of swarm systems because the swarm is assumed to be fully observable. However, due to noisy or contested environments and the assumption that the swarm is made up of low-cost agents, it is reasonable to assume that the swarm will be only partially observable to the human operator. Thus, before issuing commands to the swarm, the operator will need to spend some amount of observation time in order to approximate the state of the swarm based on samples of information from a subset of the agents.

To accurately estimate the group heading and group centroid of a flock we need a way to sample local measurements over time and then aggregate these into a global estimate. However, because we are estimating time-varying quantities, sampling for a longer period of time will not necessarily improve the accuracy of an estimate because older samples will no longer represent good estimates of the current location of the centroid.

We deal with this problem by estimating the parameters of a linear equation for the group centroid c_g as a function of time. For short periods of time, we can safely assume that the centroid is moving along a relatively linear trajectory and we can write an equation for $c_g(t)$ as

$$c_g(t) = v_g \cdot t + c_0 \quad (15)$$

where $v_g = s_g[\cos \theta_g, \sin \theta_g]^T$ is the group velocity vector, θ_g is the group heading, s_g is the group speed, and c_0 is the starting location of the swarm. We can then use least squares to find estimates \hat{v}_g and \hat{c}_0 of the parameters v_g and c_0 , respectively. Once we have obtained parameter estimates \hat{v}_g and \hat{c}_0 we can calculate an estimate of the group heading

$$\hat{\theta}_g = \text{atan2}(\hat{v}_g^y, \hat{v}_g^x). \quad (16)$$

We can additionally estimate the group centroid at some time in the near future using Equation (15). Because we are estimating a trajectory, obtaining successive samples provides a more accurate estimate of the centroid and heading of a flock.

Because the trajectory of a flock changes over time due to the stochastic inter-agent dynamics, we look at short 100 time step (10 second) intervals. While there is still some random movement, we found that breaking up the trajectory into these smaller intervals allowed us to use a linear approximation that closely approximated the actual trajectory of the flock.

To obtain the required data, we ran 10 flock simulations for 100 seconds each. We recorded the agent positions for each time step and used these positions as noisy estimates of the actual group centroid. These samples were used to find parameters to Equation (15) using least squares. We defined the true group heading over the time interval $[t_0, T]$ as

$$\theta_g = \text{atan2}(c_g^y(T) - c_g^y(t_0), c_g^x(T) - c_g^x(t_0)) \quad (17)$$

where c_g is the group centroid and $T = t_0 + 10$ seconds for our simulations. We then measured the error in predicted group heading and initial group centroid as

$$\text{error}_{\theta_g} = |\theta_g - \hat{\theta}_g| \quad (18)$$

and

$$\text{error}_{c_g} = \|c_g(t_0) - \hat{c}_0\|. \quad (19)$$

The observation time condition was split into an easy and hard parameter set. To change the difficulty of observation time we varied the number of observable agents. For easy observation time we chose to use 20 observable agents (20% of the swarm). For hard observation time we chose to use 2 observable agents (2% of the swarm). We found that $OT_{20}(5^\circ) \approx 0.2$ seconds and $OT_2(5^\circ) \approx 7.2$ seconds. At each time step sampled headings from all of the observable agents were collected and added to a buffer. For $N_{\text{obs}} = 2$ we used a buffer of 144 observations, corresponding to an observation time of 7.2 seconds of 2 agents broadcasting heading information at 10 Hz. For $Z_{\text{obs}} = 20$ we used a buffer of 40 observations, corresponding to an observation time of 0.2 seconds of 20 agents broadcasting heading and position information at 10 Hz. We calculated the circular moving average and standard deviation using the N_{obs} observed agent headings in the buffer (Fisher, 1995). This moving-average estimate of the heading of the swarm was displayed on the screen as a solid line emanating from the centroid of the observed agents with a dashed spread proportional to the circular standard deviation to provide a visual indicator of uncertainty (see (C) in Figure 12).

Observe that the model of observation time requires users to observe the samples from the flock for a period of time. Pilot results indicated that it takes a lot of time for a human to make an estimate of direction that is precise enough to allow real-time control to track a moving heading. The moving-average estimate, displayed as the solid line with dashed variance lines, acts as a perceptual aide that decreases the required observation time to a duration that allows control of the flock within time bounds of a one-hour study.

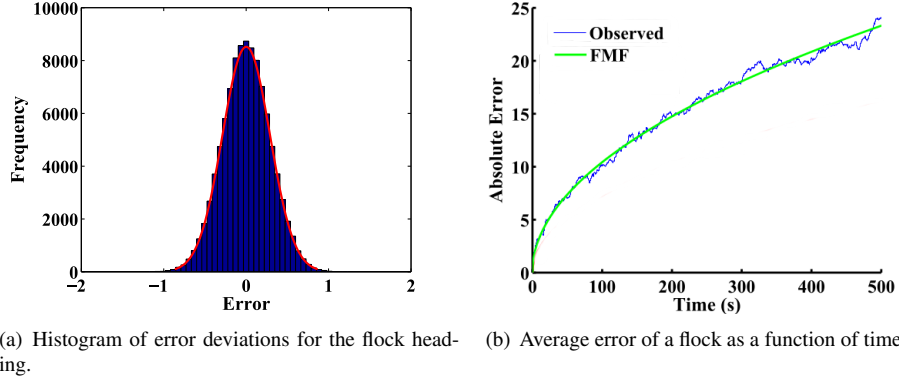


Figure 11. Flock heading natural error is approximately Gaussian and using a Wiener process results in a good fit.

5.2.3 Neglect Time Due to system noise the heading error of a flock varies stochastically with time, making it impossible to know exactly how large the error will be after t seconds. We therefore use a probabilistic measure to describe the neglect time. For our analysis, we assume that the heading error is zero at time $t = 0$. We define the probability distribution $S_{\text{flock}}(e|t)$ as the heading error distribution after t seconds. Using this distribution, we can measure the error of the flock as

$$e(t) = E[||S_{\text{flock}}(e|t)||], \quad (20)$$

the expected absolute error at time t . To model the error distribution, $S(e|t)$, we studied the data from one hundred 500-second flock simulations of the natural drift of a flock. We found that the error in the heading of a flock was approximately Gaussian, with zero-mean, uncorrelated random variables with equal variance (see Figure 11 (a)). Thus, we modeled the flock error as a one dimensional Wiener process of the form

$$S_{\text{flock}}(e|t) = \frac{1}{\sqrt{2\pi t \sigma_{\text{flock}}^2}} \exp\left(-\frac{e^2}{2t \sigma_{\text{flock}}^2}\right). \quad (21)$$

Using the average error measure in Equation (20), we can solve for the neglect time as

$$NT_{\text{flock}} = \inf_t (E[\sqrt{S_{\text{flock}}^2}] \geq e_{\text{thresh}}) \quad (22)$$

$$= \inf_t \left(\sqrt{\frac{2}{\pi} (t \sigma_{\text{flock}}^2)} \geq e_{\text{thresh}} \right) \quad (23)$$

$$= \frac{\pi}{2\sigma_{\text{flock}}^2} e_{\text{thresh}}^2 \quad (24)$$

where e_{thresh} is the error threshold, and where the second line uses the fact that $\sqrt{S_{\text{flock}}^2}$ is a folded normal distribution which has expected value $\sqrt{\frac{2}{\pi}} (t \sigma^2)$ when the normal distribution it is derived from has mean 0 and variance $t \sigma^2$ (Leone, Nelson, & Nottingham, 1961). We fit the Wiener process

model to the actual heading error to obtain our fitted model (FMF) that uses $\sigma_{S_{\text{flock}}}^2 = 1.707$ to obtain a very close fit to the actual absolute error over time (see Figure 11 (b)).

In many cases it may also be desirable to model the neglect time when the swarm is tracking a noisy or dynamic target heading. To do this we can simply add more noise to the system in the form of a noisy desired heading. In a heading tracking problem, the error is relative to the desired heading so we can add Brownian motion to the desired heading to, in effect, increase the magnitude of σ^2 and thereby model the neglect time for tracking a desired heading that evolves according to a second Wiener process with variance $t\sigma_{\text{drift}}^2$. Our equation for neglect time with a noisy heading now becomes

$$NT_{\text{flock}} = \frac{\pi}{2(\sigma_{S_{\text{flock}}}^2 + \sigma_{\text{drift}}^2)} e_{\text{thresh}}^2. \quad (25)$$

Because $\sigma_{S_{\text{flock}}}^2$ and e_{thresh} are fixed, given a desired value for NT_{flock} we can then solve for σ_{drift}^2 .

5.3 Experiment Parameters

To keep the number of conditions manageable, and because each scenario is not broken up into a specific observation phase and a specific interaction phase, we chose to lump them together and look at the combined interaction and observation time (IT+OT). We considered two conditions: Easy IT+OT (easy interaction time plus easy observation time) and Hard IT+OT (hard interaction time plus hard observation time). Adding the appropriate observation times and interaction times derived above, we have Easy IT+OT ≈ 5.5 seconds and Hard IT+OT ≈ 18.5 seconds.

We designed the user study parameters so that one scenario had an estimated fan-out of 4.5 and another had an estimated fan-out of 1. Using this information we can fill in a table as follows,

Table 1: Partially specified fan-out predictions

	Easy NT	Hard NT
Easy IT+OT	4.5	?
Hard IT+OT	?	1

We then solved for the additional noise σ_{drift}^2 to add to the system to get the desired fan-out values. To solve for σ_{drift} in the Easy NT condition we use the desired fan-out of 4.5 and Equation (12) to get Easy NT ≈ 19.25 . We then used the average error formula in Equation (25) with NT = 19.25 to get $\sigma_{\text{drift}} \approx 0.5771$. To achieve a predicted fan-out of 1, we followed the same procedure to obtain we have $\sigma_{\text{drift}} \approx 2.2859$ for the Hard NT condition. Using these values for σ_{drift} we calculated the off-diagonal entries in our condition matrix. This resulted in the 2×2 within-subject experiment in Table 2 with predicted fan-out values (rounded down) in parenthesis.

Table 2: Fan-out predictions for the four conditions investigated in our user study

	Easy NT: $\sigma_{\text{drift}} = 0.5771$	Hard NT: $\sigma_{\text{drift}} = 2.2859$
Easy IT+OT: $M_{\text{stk}} = 70, Z_{\text{obs}} = 20$	Scenario 1 (FO = 4)	Scenario 2 (FO = 2)
Hard IT+OT: $M_{\text{stk}} = 40, Z_{\text{obs}} = 2$	Scenario 3 (FO = 2)	Scenario 4 (FO = 1)

5.4 User Interface and Training

Each user was trained to a fixed level of competency by progressing through a series of tutorials that gave the user familiarity with the interface. The training scenarios were identical for each




Figure 12. Screen shot of user study GUI. (A) The visible agents in the swarm shown by circles with agent headings represented by straight lines. (B) The red cross is controlled by the user’s mouse and act as a waypoint for influencing the swarm. (C) The predicted heading of the entire swarm with a confidence range. (D) The desired heading of the swarm with dotted lines showing the error tolerance. The desired heading direction changes at each time-step according to a Wiener process with a specific standard deviation depending on the scenario. (E) The user can change their view to interact with four different swarms using the buttons on the left. (F) The triangles on the screen are to provide optical flow and simulate the movement of the swarm.

user. The training scenarios progressively introduced the users to the GUI and allowed the users to familiarize themselves with guiding a flock when all agents are visible and then guiding a flock when the swarm is only partially visible. To complete the training, users had to demonstrate an ability to switch between multiple swarms and accurately guide the swarms to follow dynamic headings. If a user failed to achieve the fixed level of competency, they repeated the training scenario until they were able to successfully complete the training.

A screen shot of the user interface is shown in Figure 12. On the screen each user could only see one swarm at a time. There are four buttons on the screen that allow the user to switch between swarms. The users interacted with the swarms by means of a mouse directed waypoint that could be turned on and off by the user. The user clicked once on the swarm display to turn on the controls and was then able to move the waypoint using the mouse to influence the swarm. Users could then click anywhere on the display to turn off the controls and click on the buttons on the screen to switch between swarms.

Users were told to try and control as many swarms as possible, but also informed that quality was important, and that for some scenarios they might need to focus on only a few swarms to achieve highly accurate tracking. Each swarm started with a random desired heading and the swarm’s initial heading was set to be 20° off course. This initial error value was chosen so that the random movement of the swarm was not likely to get the flock moving in the correct direction.

5.5 Metrics

We measured the actual fan-out by counting the number of swarms that had accurate headings for at least 50% of the scenario. We also kept track of the maximum accuracy for a single swarm for each

Table 3: Predicted fan-out and least squares means for actual measured fan-out and maximum tracking accuracy. All values are significant ($p < 0.005$).

	Scenario 1	Scenario 2	Scenario 3	Scenario 4
FO _{predicted}	4	2	2	1
FO _{actual}	2.93	1.72	1.48	0.51
Max Accuracy	87.66	65.09	66.45	46.73

scenario. After each scenario the user answered a NASA-TLX (Task Load Index) questionnaire which asked questions about perceived Mental Demand, Temporal Demand, Performance, Effort, and Frustration on each of the scenarios (Hart & Staveland, 1988). These responses were combined to obtain a raw overall NASA-TLX score.

It is important to note that when interacting with a robot or with a swarm, often the initial start-up cost of interaction will be different from the repeating interactions that occur later. To avoid these start-up costs affecting our results, we used the fact that each swarm started with a 20° error. Using our model of interaction effort, we then calculated how long it would take to turn a swarm 20° with 70 stakeholders and with 40 stakeholders. These values were then multiplied by four to get an approximate total start-up time of 30 seconds for 70 stakeholders and 56 seconds for 40 stakeholders. To eliminate the effects of start-up costs in our results we ignored the first 56 seconds of data from each scenario to give each user the predicted amount of time to correct the headings of all of the four swarms.

5.6 Results and Discussion

We recruited 34 participants from the campus of Brigham Young University and conducted the user study in accordance with IRB protocol. One of the participants failed to answer several of the survey questions and so we did not use that user’s data. This left us with data from 33 participants, 22 male and 11 female between the ages of 17 and 33. No adverse side effects were reported during the user study.

To analyze our results we used SAS to run a mixed-models analysis with Tukey-Kramer adjustment blocking on subjects (Tukey, 1949). We analyzed both fixed effects and pairwise within-subject differences across scenario. We found that gender was close to significant but was confounded by video game experience so we decided to keep gender as part of the analysis.

5.6.1 Fan-out The results for fan-out are shown in Table 3. The first row of data shows the values we predicted based on the analysis contained in this paper. The second row shows the actual fan-out computed by counting the number of swarms with headings within 5° of the desired heading for at least 50% of the simulation. Note that the neglect tolerance of 5° was subjectively chosen to balance between a desire for good performance and the need to respect human variability. Note further that by making the 5° criterion higher or the 50% standard lower, we get higher values for predicted or actual fan-out, respectively. The last row in Table 3 shows the average maximum accuracy for each of the four scenarios. The mixed models analysis of variance for fixed effects showed that FO_{actual} was significantly affected by scenario ($F(3,96)=34.77, p < 0.0001$) but not gender ($F(1, 96) = 1.41, p = 0.2383$) and that Max Accuracy was significantly affected by scenario ($F(3, 96) = 38.32, p < 0.0001$) but not gender ($F(1, 96) = 1.44, p = 0.2334$).

We see that while the actual fan-out values calculated from the user study are lower than predicted, the general trends from the user study closely match the predicted trend. For example, comparing the predicted fan-out to the actual fan-out, we see that scenario 1 has a much higher value

Table 4: Least squares means for raw NASA-TLX survey results and perceived difficulty for interaction, neglect, and observation. All values are significant ($p < 0.0001$).

	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Raw NASA-TLX	220.23	254.47	286.29	323.56

than the other scenarios, scenario 2 and 3 are very similar, and scenario 4 is definitely the hardest. Indeed, the statistical results on the differences of least squared means between the different scenarios showed that for FO_{actual} , the differences between all scenarios were significant ($p < 0.001$) except for scenarios 2 and 3 ($p = 0.7418$). Finally, if we examine the highest accuracy achieved on average for each scenario we see that clearly scenario 1 is the easiest, scenarios 2 and 3 are of medium difficulty with no significant difference ($p = 0.9844$), and scenario 4 is the hardest.

5.6.2 Perceived workload and difficulty The results for raw NASA-TLX are shown in Table 4. The fixed effects analysis showed that scenario significantly influenced Raw NASA-TLX ($F(3, 96) = 27.57$, $p < 0.0001$). Gender did not significantly influence Raw NASA-TLX ($F(1, 96) = 1.89$, $p = 0.1730$).

Examining the raw NASA-TLX scores we see an ascending level of perceived workload with scenario 1 the easiest and scenario 4 the hardest. The differences in least squares means were all significant ($p < 0.05$). The least significant difference was again between scenarios 2 and 3 ($p = 0.0428$). These scenarios were designed to be almost identical in terms of fan-out, but we see that in terms of perceived workload they are significantly different.

5.7 Discussion

The motivation for the user study was that if a flock is single perceptual unit then a human should be able to manage multiple flocks in a predictable way. This was operationalized into a hypothesis that it should be possible to predict human workload in the form of fan-out for maintaining specific time-varying target headings. The above results show that the individual scenarios that we designed had statistically significant impact on every metric examined. While the actual fan-out values calculated from the user study were lower than predicted, the general trends from the user study roughly match the predicted trends based on our modeling and analysis. Despite the low accuracy for predicting actual fan-out values, the results do demonstrate that users were able to manage multiple flocks as collective units. This provides additional evidence that collective state is a fundamental percept that allows a human to simultaneously manage the behaviors of multiple collectives through interaction, observation, and neglect. Future work will investigate how to improve the accuracy of fan-out predictions.

In contrast to the results on fan-out, the NASA-TLX scores for scenarios 2 and 3 did not follow our predicted trend. Table 4 shows that scenario 3 was perceived to be significantly harder than scenario 2, despite very similar actual accuracy and performance over both tasks (see Table 3). Thus, despite similar accuracy and performance, users felt they had to work harder and concentrate more when using fewer agents to solve a simpler tracking problem (scenario 3), than when using more agents to solve a difficult tracking problem (scenario 2). We discuss how these results relate to our second proposed invariant, span and persistence, in the next section.

We also note that as a side effect of the study, the experiment results show that human-swarm interactions for spatially coherent flocks can be approximately described using three system characteristics: interaction time, neglect time, and observation time. The system characteristics triple, (IT, NT, OT) , provides valuable system-level information on human-swarm interactions that have

the potential to allow researchers and practitioners the ability to estimate human workload and performance in a variety of situations. Additionally, these system characteristics provide information about the maximum number of flocking agents that a human can control. As noted previously, one of the limitations of our study is the assumption that IT and OT are separable. Future work should examine more realistic methods of modeling IT and OT that take into account that interactions and observations are often interleaved. Better modeling of these quantities will likely provide more accurate estimates of fan-out and workload.

6. Invariant Two: Span and Persistence

The previous sections provide evidence that a human can perceive and influence a spatially coherent swarm as a collective unit. The evidence for supporting this was mostly indirect, meaning that we never directly measured how humans perceived a swarm as opposed to, say, a decoherent set of randomly-moving independent agents. As mentioned in the section on related literature, recent results have made attempts to directly measure how a human perceives a swarm; initial results, though encouraging, still need further validation (Manning et al., 2015; Harriott et al., 2014). Trends in the recent work provide further support for the claim of perception of collective unit.

Importantly, perception of a swarm as a collective unit is not the only invariant of how humans manage the spatio-temporal swarms discussed in this paper. A second invariant is what we call “span and persistence”; *span* means the number of agents that the human manages and *persistence* means how long the human influences those agents. We claim that *there is a fundamental trade-off between span and persistence in human-swarm interaction*. Simply put, if you can control all the agents at once you can make the group change quickly; however, if you can only control one agent in the group, then you may either have to persist a long time to effect change, or it may be impossible to effect that change at all.

Most of the results in the previous sections touch on this trade-off between span and persistence. This section, though brief, discusses how the trade-off manifests itself. Rather than repeating the linear order of the experiments from the previous section, we instead take a more conceptual approach, organizing the discussion around common forms of influence or perception.

Consider the studies in (Pendleton & Goodrich, 2013; Goodrich et al., 2012), where human influence occurred through a single human-controlled leader or human-controlled predator. The problem with the predator, which influenced the agents by repelling them, was that the human could not easily persist in exerting influence. Because the agents “ran away” from the predator, moving the predator only allowed the human to influence any individual or small group of agents for a brief period of time. Large spans of influence were not sustainable and long periods of persistence were difficult to obtain. The result from both the user studies and the simulation studies is that exerting human influence via a predator took a lot of human effort and tended to cause a swarm to fragment.

Exerting influence through a leader, by contrast, allowed a human to effectively control a flock and to do so with low workload. The problem was that a leader agent tended to exert too much influence. As agents came in contact with the leader, the attraction to the leader kept the agent within the sphere of influence of the leader. Span grew over time and persistence grew too because agents wouldn’t voluntarily leave the sphere of influence. This was fine for leading flocks as long as the leader was moving, but using this style of leadership for a torus was impossible because the torus eventually decayed into a flock with all agents following the leader.

The “sweet spot” was a balance between the predator’s low span and persistence, and the leader’s high span and persistence. A set of stakeholders was able to strike this balance well, and at this balance point the trade-off between span and persistence emerged. The first place we saw this was in Figures 3 and 4. These figures contained two axes, an explicit span axis, the number of agents in the collective who served as stakeholders, and an implicit persistence axis, the priority parameters.

When enough agents were assigned as stakeholders, it was fairly easy to switch behaviors. When priority was too low, the goal point was overwhelmed by the influence of nominal agents, resulting in a form of persistence where human input wasn't important enough to draw agents to the goal.

It is interesting that the style of the stakeholder in Figures 3 and 4 was important. Leadership by orientation was more effective at turning a flock into a torus because it tended to cause the stakeholders to align with the human's intent. Like the priority parameter from the previous section, this creates a "trickle-down" form of human influence where agents that tend to align with human intent for longer have more influence over the collective. The basic characteristic of a flock is that all agents are oriented the same way, so human intent persists with the flock better when leadership emphasizes orientation. A similar argument holds for having humans exert influence through attraction so that their influence "trickles down" more persistently.

We saw something similar in (Kerman et al., 2012) and (Kerman, 2013) where there was a trade-off between how many agents the human could influence and how quickly the torus or flock could track a moving point. Results indicated something that is unsurprising given the understanding of span and persistence. If the stakeholders in a torus moved too quickly in order to track a moving point, they either broke away from the torus or they caused the torus to turn into a flock—examples of too little and too much persistence, respectively (Brown, Jung, & Goodrich, 2014).

Trade-offs in span and persistence appear not only in exerting influence over a swarm, but also appear in perceiving the behavior of the swarm. The sample trajectories shown in figures 5 and 6 give an example of how watching a few agents for a long period of time gives a similar confidence in the collective behavior as does watching many agents for a short period of time. As a concrete example, the classifier we used to determine whether a swarm was in flock or swarm state, which we published in (Brown & Goodrich, 2014), had explicit trade-offs between span and persistence. In this case, persistence means the amount of time required to recognize whether the swarm was in flock or torus state. As shown in Figure 7, if the classifier had access to many agents then classification went quickly, but if the classifier had access to only a few agents then classification went slowly.

Finally, the user study from the previous section had a span and persistence trade-off. In this case, rather than trying to classify whether the swarm was in a flock or torus state, the objective was to determine the direction that a flock was heading, and interact with the flock in order to correct its heading. We fixed span in the user study, but persistence was left to the user. We found that as both the interaction span and observation span were reduced, the participants in the study were unable to persistently influence each of the four flocks to maintain proper headings. This was manifest in the results in Table 3 that showed a dramatic decrease in fan-out when moving from high span (Scenario 1) to low span (Scenario 4). We also see the trade-off between span and persistence in the results of Table 4. The NASA-TLX scores show that the perceived workload increased as the span decreased and as the noise in the scenarios increased. Participants felt that they had to work harder (persistently interact more frequently and for longer periods of time) when span was reduced. Thus, higher span requires lower persistence and results in less operator load and higher fan-out, while lower span requires higher persistence and results in higher operator workload and lower fan-out.

In summary, the above evidence strongly supports the claim that span and persistence impact all aspects of Human-Swarm Interaction that we have examined in this paper. Leading a flock or torus through stakeholders requires balancing the trade-off between span and persistence; however, influence through a leader or predator doesn't properly balance this trade-off. Similarly, switching between a flock and a torus, classifying whether a swarm is in a flock or torus state, and managing the performance of multiple flocks all have trade-offs between span and persistence that impact the effective and successful interactions between a human and a bio-inspired swarm.

7. Conclusions and Future Work

While the search for invariants is as old as science itself, the search for invariants in Human-Swarm Interaction is still in its infancy. The aim of this paper is to provide a first attempt at understanding and formalizing some of the invariants of HSI. In the preceding sections we have synthesized results from mathematical analysis, simulations, and user-studies to provide evidence of two fundamental invariants underlying many facets of HSI: (1) *collective state* is the fundamental percept associated with a bio-inspired swarm and (2) a human’s ability to influence and understand the collective state of a swarm is determined by the fundamental balance between *span and persistence*. While we note that these invariants are by no means complete or exhaustive, we believe that these invariants will provide researchers in the field of HSI with key insights and will benefit future studies by enabling useful abstraction, prediction, analysis, and modeling.

The proposed invariant of collective state implies that human interactions with bio-inspired swarms should focus on the higher-level attractor-states of the collective. Swarm behaviors that are attractors are more amenable to human interaction and enable collective-level switches in behavior, while also limiting the required human influence and bandwidth required for these interactions to be effective. We propose that focusing on managing the stable states of a collective, rather than managing individuals in the swarm, provides an ideal way allow a single human to manage hundreds of robots simultaneously. However, there are still many gaps in understanding how humans can manage swarms and other multi-agent systems at the attractor level. Future work should seek to better understand how local perturbations can shift the collective state of a multi-agent system between basins of attraction, causing global state changes. Research should also investigate what kinds of interfaces are most useful for controlling and influencing swarms with limited interactions. In particular there, future work should consider additional ways in which an operator can manage multiple swarms, as was done in Section 5. Changing the waypoint input to a simple directional heading input would be an obvious improvement.

Our research has focused on geometric properties of swarms resulting from simple attraction, repulsion, and orientation behaviors. We hypothesize that these invariants also hold for other swarm and colony behaviors that can be described by dynamical systems; however, there are still many open questions. How generalizable are these concepts if there is an actual goal for the robots other than modification of their geometric properties? Is it important to be able to observe or control other aspects of state, and what repercussions do the invariants have on the kinds of information that need to be observed or the limitations on input controls? We believe these are important questions that should be addressed in future work.

The proposed invariant of span and persistence provides an important key to understanding how control can be appropriately shared between a human operator and a bio-inspired swarm. Often a human will have knowledge of global objectives and high-level plans, but members of the collective will often have more detailed low-level knowledge about environmental particulars. In HSI, a human’s ability to successfully interact with a swarm is fundamentally tied to the span and persistence of their interactions with the swarm. Too much span or persistence can cause the human to override the inherent stability and collective properties of a swarm; however, too little span or persistence can preclude the human from ever influencing the swarm’s behavior.

Preliminary work on formalizing the theory of span and persistence has produced inconclusive results (Brown, Jung, & Goodrich, 2014). Thus, we feel that this is an important area for further research with many open questions. The invariant of span and persistence has many implications for future HSI. How to design interfaces that allow an operator to visualize span and persistence and find the optimal balance between them remains an intriguing and challenging problem. We also note the interesting connection between the mechanism of quorum sensing found in many biological systems (Seely, 2010) and the proposed invariant of span and persistence. For example, the

collective decision made by house-hunting honeybees is directly related to the number (span) and persistence of the individual scouts. In previous work, we have shown that quorum sensing provides a mechanism for changing the required span and persistence needed to change the behavior of a swarm (Brown, Kerman, & Goodrich, 2014). This allows a human to find the appropriate balance between the controllability of the swarm and the vulnerability of the swarm to agent failures and perturbations. We hypothesize that the invariants of collective state and of span and persistence can be seen in many forms of collective intelligence, especially in the context of humans interacting with higher-order collectives, such as bio-inspired colonies and packs. We are currently investigating this hypothesis as it relates to honeybee-inspired collectives.

Finally, future work should also examine the concept of neglect benevolence (Nagavalli et al., 2015) as another possible invariant of HSI. Neglect benevolence is related to our first invariant because collective states that are attractors of a complex dynamical system allow a human to neglect a swarm between interactions. This can often be beneficial because it allows transient behaviors to die off, allowing the swarm to fully stabilize before new human influence is inserted into the swarm. We also hypothesize that neglect benevolence is related to, and possibly complements our proposed invariant of span and persistence.

Acknowledgements

We would like to thank the Office of Naval Research Science of Autonomy Program and the Air Force Information Directorate Machine Intelligence for Mission Focused Autonomy Program for funding this research. The findings, opinions, and conclusions of this paper do not necessarily reflect the views of the funding agencies. We would also like to thank Dennis Eggett of the BYU statistical consulting center for help with the statistical analysis and the anonymous reviewers for their many insightful and helpful comments and suggestions.

References

- Alboul, L., Saez-Pons, J., & Penders, J. (2008, October). Mixed human-robot team navigation in the GUARDIANS project. In *Proceedings of the international conference on safety, security, and rescue robotics*. Sendai, Japan.
- Baddeley, A. (1986). *Working memory* (Vol. 11). Oxford.
- Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., et al. (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences*, 105(4), 1232–1237.
- Bashyal, S., & Benayagamoorthy, G. K. (2008). Human swarm interaction for radiation source search and localization. In *Ieee swarm intelligence symposium*. St. Louis, Missouri, USA.
- Bashyal, S., & Venayagamoorthy, G. K. (2008). Human swarm interaction for radiation source search and localization. In *Swarm intelligence symposium* (pp. 1–8).
- Bode, N., Franks, D., & Wood, A. (2010). Limited interactions in flocks: relating model simulations to empirical data. *Journal of The Royal Society Interface*, 8(55), 301–304.
- Brown, D. S., & Goodrich, M. A. (2014). Limited bandwidth recognition of collective behaviors in bio-inspired swarms. In *Proceedings of the 2014 international conference on autonomous agents and multi-agent systems* (pp. 405–412).
- Brown, D. S., Jung, S.-Y., & Goodrich, M. A. (2014). Balancing human and inter-agent influences for shared control of bio-inspired collectives. In *Proceedings of IEEE international conference on systems, man, and cybernetics*. San Diego, California.
- Brown, D. S., Kerman, S. C., & Goodrich, M. A. (2014). Human-swarm interactions based on managing attractors. In *Proceedings of the 2014 ACM/IEEE international conference on human-robot interaction* (pp. 90–97).
- Chitsaz, H., & LaValle, S. M. (2007). Time-optimal paths for a dubins airplane. In *Decision and control, 2007 46th IEEE conference on* (pp. 2379–2384).

- Conradt, L., Krause, J., Couzin, I., & Roper, T. (2009). Leading according to need in self-organizing groups. *The American Naturalist*, 173(3), 304–312.
- Couzin, I., & Krause, J. (2003). Self-organization and collective behavior in vertebrates. *Advances in the Study of Behavior*, 32, 1–75.
- Couzin, I., Krause, J., Franks, N., & Levin, S. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025), 513–516.
- Couzin, I., Krause, J., James, R., Ruxton, G., & Franks, N. (2002). Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1), 1–11.
- Crandall, J., Goodrich, M., Olsen Jr, D., & Nielsen, C. (2005). Validating human-robot interaction schemes in multitasking environments. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 35(4), 438–449.
- Cummings, M. L. (2004). Human supervisory control of swarming networks. In *2nd annual swarming: Autonomous intelligent networked systems conference* (pp. 1–9).
- Ding, X. C., Rahmani, A. R., & Egerstedt, M. (2010). Multi-uav convoy protection: an optimal approach to path planning and coordination. *Robotics, IEEE Transactions on*, 26(2), 256–268.
- Dubins, L. E. (1957). On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American J. of Mathematics*, vol. 79, no. 3.
- Feynman, R. P. (1964). *The character of physical law* (1994Edition ed.). Modern Library.
- Fisher, N. I. (1995). *Statistical analysis of circular data*. Cambridge University Press.
- Genter, K., Agmon, N., & Stone, P. (2013). Ad hoc teamwork for leading a flock. In *Proceedings of the 2013 international conference on autonomous agents and multi-agent systems* (pp. 531–538).
- Glas, D. F., Kanda, T., Ishiguro, H., & Hagita, N. (2012). Teleoperation of multiple social robots. In (Vol. 42, pp. 530–544). IEEE.
- Goodrich, M. A. (2010). On maximizing fan-out: Towards controlling multiple unmanned vehicles. In M. Barnes & F. Jentsch (Eds.), *Human-robot interactions in future military operations*. Surrey, England: Ashgate Publishing.
- Goodrich, M. A., & Mercer, E. G. (2012, June). Abstraction and persistence: Macro-level guarantees of collective bio-inspired teams under human supervision. In *Proceedings of Infotech@Aerospace*. Garden Grove, California.
- Goodrich, M. A., Pendleton, B., Kerman, S., & Sujit, P. B. (2012, June). What types of interactions to bio-inspired robot swarms and flocks afford a human? In *Proceedings of robotics science and systems*. Sydney, Australia.
- Harriott, C. E., Seiffert, A. E., Hayes, S. T., & Adams, J. A. (2014). Biologically-inspired human-swarm interaction metrics. In *Proceedings of the human factors and ergonomics society annual meeting* (Vol. 58, pp. 1471–1475).
- Hart, S. G., & Staveland, L. E. (1988). Development of nasa-tlx (task load index): Results of empirical and theoretical research. *Advances in psychology*, 52, 139–183.
- Herbert-Read, J., Perna, A., Mann, R., Schaerf, T., Sumpter, D., & Ward, A. (2011). Inferring the rules of interaction of shoaling fish. *Proceedings of the National Academy of Sciences*, 108(46), 18726–18731.
- Jadbabaie, A., Lin, J., & Morse, A. (2003). Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6), 988–1001.
- Jung, S.-Y., Brown, D. S., & Goodrich, M. A. (2013). Shaping couzin-like torus swarms through coordinated mediation. In *Systems, man, and cybernetics (smc), 2013 ieee international conference on* (pp. 1834–1839).
- Jung, S.-Y., & Goodrich, M. A. (2013). Multi-robot perimeter-shaping through mediator-based swarm control. In *Proceedings of the 16th International Conference on Advanced Robotics (ICAR)* (pp. 1–6).
- Kerman, S. C. (2013). *Methods and metrics for human interaction with bio-inspired robot swarms*. Unpublished master's thesis, Brigham Young University.
- Kerman, S. C., Brown, D. S., & Goodrich, M. A. (2012). Supporting human interaction with robust robot swarms. In *Proceedings of the international symposium on resilient control systems*.
- Kira, Z., & Potter, M. (2010). Exerting human control over decentralized robot swarms. In *4th international conference on autonomous robots and agents* (pp. 566–571).

- Kolling, A., Sycara, K., Nunnally, S., & Lewis, M. (2013). Human swarm interaction: An experimental study of two types of interaction with foraging swarms. *Journal of Human-Robot Interaction*, 2(2), 103–128.
- Lee, J. D., & See, K. A. (2003). Trust in automation: Designing for appropriate reliance. *Human Factors*. (To appear.)
- Leone, F., Nelson, L., & Nottingham, R. (1961). The folded normal distribution. *Technometrics*, 3(4), 543–550.
- Levine, H., Rappel, W., & Cohen, I. (2000). Self-organization in systems of self-propelled particles. *Physical Review E*, 63(1), 017101.
- Lien, J.-M., Bayazit, O. B., Sowell, R. T., Rodriguez, S., & Amato, N. M. (2004). Shepherding behaviors. In *Proceedings of the IEEE international conference on robotics and automation*. New Orleans, Louisiana.
- Manning, M. D., Harriott, C. E., Hayes, S. T., Adams, J. A., & Seiffert, A. E. (2015). Heuristic evaluation of swarm metrics' effectiveness. In *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction extended abstracts* (pp. 17–18). New York, NY, USA: ACM. Available from <http://doi.acm.org/10.1145/2701973.2702046>
- Marshall, J., Broucke, M., & Francis, B. (2004). Formations of vehicles in cyclic pursuit. *IEEE Transactions on Automatic Control*, 49(11), 1963–1974.
- Mesbahi, M., & Egerstedt, M. (2010). *Graph theoretic methods in multiagent networks*. Princeton University Press.
- Nagavalli, S., Chien, S.-Y., Lewis, M., Chakraborty, N., & Katia, S. (2015). Bounds of neglect benevolence in input timing for human interaction with robotic swarms. In *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction* (pp. 197–204).
- Nagavallik, S., Luo, L., Chakraborty, N., & Sycara, K. (2014). Neglect benevolence in human control of robotic swarms. In *IEEE international conference on robotics and automation (icra)* (pp. 6047–6053).
- Nevai, A. L., & Passino, K. M. (2010). Stability of choice in the honey bee nest-site selection process. *Journal of Theoretical Biology*, 263(1), 93–107.
- Nunnally, S., Walker, P., Lewis, M., Kolling, A., Chakraborty, N., & Sycara, K. (2012). Robotic swarm connectivity with human operation and bandwidth limitations. In *Proceedings of the AAAI fall symposium workshop on human control of bio-inspired swarms*.
- Olfati-Saber, R. (2006). Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3), 401–420.
- Olsen Jr, D. R., & Wood, S. B. (2004). Fan-out: Measuring human control of multiple robots. *Proceedings of the SIGCHI conference on Human factors in computing systems*, 231–238.
- Olsen Jr, D. R., Wood, S. B., & Turner, J. (2004). Metrics for human driving of multiple robots. *2004 IEEE International Conference on Robotics and Automation*, 3, 2315–2320.
- Pashler, H. (1998). *The psychology of attention*. Cambridge, Massachusetts, USA: MIT Press.
- Pendleton, B., & Goodrich, M. (2013). Scalable human interaction with robotic swarms. *Proceedings of the AIAA Infotech@Aerospace Conference*.
- Pourmehr, S., Monajjemi, V. M., Vaughan, R., & Mori, G. (2013). "you two! take off!": Creating, modifying and commanding groups of robots using face engagement and indirect speech in voice commands. In (pp. 137–142).
- Ren, W., & Beard, R. (2007). *Distributed consensus in multi-vehicle cooperative control: theory and applications*. Springer.
- Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioral model. In *ACM SIGGRAPH computer graphics* (Vol. 21, pp. 25–34).
- Romero, L., Forgoston, E., & Schwartz, I. (2011). Noise, bifurcations, and modeling of interacting particle systems. In *International conference on intelligent robots and systems* (pp. 3905–3910).
- Rubenstein, M., Cabrera, A., Werfel, J., Habibi, G., McLurkin, J., & Nagpal, R. (2013). Collective transport of complex objects by simple robots: theory and experiments. In *Proceedings of the 2013 international conference on autonomous agents and multi-agent systems* (pp. 47–54).
- Seely, T. D. (2010). *Honeybee democracy*. Oxfordshire, UK: Princeton University Press.
- Seiffert, A. E., Hayes, S. T., Harriott, C. E., & Adams, J. A. (2015). Motion perception of biological swarms. In *Proceedings of the annual cognitive science society meeting*.
- Sheridan, T. B. (2002). *Humans and automation: System design and research issues*. John Wiley and Sons.

- Simon, H. A. (1990). Invariants of human behavior. *Annu. Rev. Psychology*, 41, 1-19.
- Spears, W., Spears, D., Heil, R., Kerr, W., & Hettiarachchi, S. (2005). An overview of physicomimetics. *Swarm Robotics*, 84-97.
- Steinfeld, A., Jenkins, O. C., & Scassellati, B. (2009). The oz of wizard: simulating the human for interaction research. In *Human-robot interaction (hri), 2009 4th acm/ieee international conference on* (pp. 101-107).
- Stilwell, D. J., & Bay, J. S. (1993). Toward the development of a material transport system using swarms of ant-like robots. In *Proceedings of the ieee international conference on robotics and automation* (pp. 766-771).
- Strömbom, D. (2011). Collective motion from local attraction. *Journal of theoretical biology*, 283(1), 145-151.
- Su, H., Wang, X., & Lin, Z. (2009). Flocking of multi-agents with a virtual leader. *IEEE Transactions on Automatic Control*, 54(2), 293-307.
- Sumpter, D. J. T. (2010). *Collective animal behavior*. Princeton University Press.
- Tukey, J. W. (1949). Comparing individual means in the analysis of variance. *Biometrics*, 99-114.
- Vaughn, R., Sumpter, N., Henderson, J., Frost, A., & Cameron, S. (1998). Robot control of animal flocks. In *Proceedings of the international symposium of intelligent control* (p. 277-282). Gaithersburg, Maryland.
- Vicsek, T., Czirók, A., Ben-Jacob, E., Cohen, I., & Shochet, O. (1995). Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6), 1226-1229.
- Walker, P., Nunnally, S., Lewis, M., Kolling, A., Chakraborty, N., & Sycara, K. (2012). Neglect benevolence in human control of swarms in the presence of latency. In *Proceedings of the IEEE international conference on systems, man, and cybernetics* (pp. 3009-3014).
- Werfel, J., Petersen, K., & Nagpal, R. (2014). Designing collective behavior in a termite-inspired robot construction team. *Science*, 343(6172), 754-758.
- Wickens, C. D., & Hollands, J. G. (2000). *Engineering psychology and human performance* (Third ed.). Prentice-Hall, New Jersey.
- Zheng, K., Glas, D. F., Kanda, T., Ishiguro, H., & Hagita, N. (2013a). Designing and implementing a human-robot team for social interactions. In (Vol. 43, p. 843-859). IEEE.
- Zheng, K., Glas, D. F., Kanda, T., Ishiguro, H., & Hagita, N. (2013b). Supervisory control of multiple social robots for navigation. In *Proceedings of the 8th acm/ieee international conference on human-robot interaction* (p. 17-24).

Authors' contact information: Daniel S. Brown, AFRL Information Directorate, Rome NY, USA. Email: daniel.brown.81@us.af.mil. Michael A. Goodrich, Brigham Young University, Provo, UT USA. Email: mike@cs.byu.edu. Shin-Young Jung, Brigham Young University. Email: ldsrogan@gmail.com. Sean Kerman, Brigham Young University. Email: sean.kerman@gmail.com.